

Please type a plus sign (+) inside this box [+]

PTO/SB/05 (12/97)

Approved for use through 09/30/00. OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number



UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Attorney Docket No. 042390.P7279

Total Pages 2

First Named Inventor or Application Identifier Nimrod Diamant

Express Mail Label No. EL414970377US

ADDRESS TO: Assistant Commissioner for Patents
Box Patent Application
Washington, D. C. 20231

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

1. X Fee Transmittal Form
(Submit an original, and a duplicate for fee processing)
2. X Specification (Total Pages 26)
(preferred arrangement set forth below)
 - Descriptive Title of the Invention
 - Cross References to Related Applications
 - Statement Regarding Fed sponsored R & D
 - Reference to Microfiche Appendix
 - Background of the Invention
 - Brief Summary of the Invention
 - Brief Description of the Drawings (if filed)
 - Detailed Description
 - Claims
 - Abstract of the Disclosure
3. X Drawings(s) (35 USC 113) (Total Sheets 8)
4. X Oath or Declaration (Total Pages 4)
 - a. X Unsigned
 - b. Copy from a Prior Application (37 CFR 1.63(d))
(for Continuation/Divisional with Box 17 completed) (**Note Box 5 below**)
 - i. DELETIONS OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b).
5. Incorporation By Reference (useable if Box 4b is checked)
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. Microfiche Computer Program (Appendix)

1520 U.S. 11/18/99



7. ☒ Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)
a. ☐ Computer Readable Copy
b. ☐ Paper Copy (identical to computer copy)
c. ☐ Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

8. ☐ Assignment Papers (cover sheet & documents(s))
9. ☐ a. 37 CFR 3.73(b) Statement (where there is an assignee)
☒ b. Power of Attorney
10. ☐ English Translation Document (if applicable)
11. ☐ a. Information Disclosure Statement (IDS)/PTO-1449
☐ b. Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Return Receipt Postcard (MPEP 503) (Should be specifically itemized)
14. ☐ a. Small Entity Statement(s)
☐ b. Statement filed in prior application, Status still proper and desired
15. ☐ Certified Copy of Priority Document(s) (if foreign priority is claimed)
16. ☒ Other: Certificate of Mailing (1 page in duplicate)

17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information:
☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP)
of prior application No:

18. **Correspondence Address**
☐ Customer Number or Bar Code Label
(Insert Customer No. or Attach Bar Code Label here)
or
☒ Correspondence Address Below

NAME Steven D. Yates, Reg. No 42,242
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

ADDRESS 12400 Wilshire Boulevard
Seventh Floor

CITY Los Angeles STATE California ZIP CODE 90025-1026
Country U.S.A. TELEPHONE (503) 684-6200 FAX (503) 684-3245

Express Mail Label: EL414970377US

12/01/97

-2-

PTO/SB/05 (12/97)

Approved for use through 09/30/00. OMB 0651-0032
Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

65617-920000

UNITED STATES PATENT APPLICATION
FOR
**Apparatus and Method for
Networking Driver Protocol Enhancement**

INVENTOR:

Nimrod Diamant
4 Haegoz St.
Kfar-Saba, Israel 44418

Prepared by

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN, LLP
12400 WILSHIRE BOULEVARD
SEVENTH FLOOR
LOS ANGELES, CALIFORNIA 90025-1026

(503) 684-6200

Express Mail mailing label number: EL 41497037745

Date of Deposit: 11/18/99

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231.

Heather L. Adamson
Typed or printed name of person mailing paper or fee

[Signature]
Signature of person mailing paper or fee

11.18.99
Date signed

06677-9206760

Apparatus and Method for Networking Driver Protocol Enhancement

Field of the Invention

5 The invention generally relates to networking, and more particularly to runtime manipulation of networking drivers.

Background

10 Since the advent of computers, various techniques have been used to allow computers to communicate. The common term for such communication is “networking” and many different models have been developed to describe the communication process. A well-known approach to representing networking communication is through use of the Open Systems Interconnection (OSI) Reference Model, which is a seven-layer model put forth by the International Standards Organization (ISO) in 1983.

15 The seven layers are the physical, data link, network, transport, session, presentation, and application layers; each layer represents a successive level of complexity and data-abstraction for an underlying raw data stream. That is, the physical layer (part of a network interface card (NIC)) is generally responsible for transmitting and receiving raw data (bits) from a network medium or interface. The physical layer
20 does not maintain or require any particular structure to the data. The data link layer, however, abstracts the raw data stream into a series of frames, or chunks of raw data, allowing this layer to be generally responsible for ensuring that data is correct (e.g., all data received is correct). The network layer concerns properly routing data to and from appropriate sources and destinations. Each remaining layer is similarly responsible for
25 providing further abstraction and integrity to a data stream. (For further information regarding the OSI model, see *Computer Networks* by Andrew Tanenbaum, Prentice Hall (2d. Ed. 1989).)

In order to distinguish one NIC from another, the NICs each have a unique identifier, commonly referred to as a Media Access Control (MAC) identifier.

30 Communication protocols are bound to one or more NICs, allowing data utilizing the protocol to be routed over bound NICs. Commonly used protocols include connection-oriented protocols such as Transmission Control Protocol (TCP) and Sequenced Packet

Exchange (SPX), as well as connectionless protocols such as Internet Protocol (IP), Internetwork Packet Exchange (IPX), and User Datagram Protocol (UDP).

Binding a NIC with a protocol is normally an operating-system boot-time event. For networking, after a machine performs its power-on self test (POST), and the operating system begins to load, control is passed to networking drivers which bind an instance of themselves to each NIC, and to one or more protocol stacks. The drivers also request its associated NIC to provide the NIC's MAC address. This address is stored in a memory reserved for the driver, and is also stored in the NIC's receive address filtering hardware. Once all NIC drivers have been initialized, data can be routed to and from each NIC.

Unfortunately, as will become clear in the following detailed description, once the Protocol stacks are bound to NICs using their drivers, there is no way to handle a NIC failure. If a particular NIC is used in a communication session between a source and a destination computer, and the NIC fails, current networking software does not allow silent replacement of the failed NIC with an operative NIC. Instead, a communication error occurs, preventing the application programs from being able to continue communication.

Some attempts have been made to overcome NIC failure issue. One solution is to overcome NIC failure by using standard functions to identify routing changes. If a NIC changes, a broadcast of a routing change is made to all clients, indicating that they need to switch to a new destination address (e.g., the replacement NIC). This solution, however, lacks transparency. Another alternate, as discussed in U.S. Patent No. 5,661,719, is to provide NICs with programmable MAC addresses. This allows a new NIC to emulate the failed NIC, but it is a solution lacking flexibility.

Summary

One aspect of the invention is an application programming interface (API) for enhancing data network communication. The API includes an identify address function including programming instructions for identifying a stored node address stored by a base driver for a network interface associated with the base driver. The API also includes an update node address function including programming instructions for

directing the base driver to update the stored node address with a new node address in a configuration storage of the base driver, and in a receive address filtering table for the network interface.

5

Brief Description of the Drawings

Features and advantages of the invention will become apparent to one skilled in the art to which the invention pertains from review of the following detailed description and claimed embodiments of the invention, in conjunction with the drawings in which:

FIG. 1 illustrates a network communication configuration.

10 FIG. 2 shows the relationship between the traditional OSI model and Novell's ODI API.

FIG. 3 illustrates one implementation of the FIG. 1 communication configuration in the FIG. 2 Novell networking context.

FIG. 4 is a general flowchart for transparent fail-over for a defective NIC.

15 FIG. 5 is a general flowchart for transparently re-packaging an original data stream in a new protocol format.

FIGS. 6 shows API function entries for implementing on-the-fly modifications to networking communication.

20 FIG. 7 illustrates a suitable computing environment in which the claimed invention may be practiced.

Detailed Description

FIG. 1 illustrates a network communication configuration, in which a protocol stack **100** is in communication with an intermediary layer **102** (e.g., LSL or NDIS).

25 There may, as illustrated, be several protocol stacks **100**. It is assumed there is only a single protocol stack and a single intermediary layer. The protocol stack corresponds to typical networking protocols such as TCP, IP, SPX, IPX, NetBios, Netbeui, AppleTalk, X.400, and the like. The intermediary layer **102** is bound to the protocol stack, and helps route network traffic. The intermediary layer is in communication with multiple
30 network interface card base drivers **104-108**. As shown, instances of a single base driver **104** can be managing multiple network interfaces (three such interfaces are illustrated as a stack of interfaces **116**). For presentation clarity, it is assumed each

base driver communicates with a single network interface. Note that although network interface cards, or "NICs", are shown, the term NIC is meant to include input/output interfaces for alternate network configurations, such networks effected over serial/parallel port connections, Universal Serial Bus (USB) links, IEEE 1394 FireWire link, and the like.

In the illustrated configuration, the intermediary **102** appears to the stack **100** as a multiplexer to the different base drivers. The stack and base drivers are bound to the intermediary, resulting in network data received by the protocol stack being routed to the intermediary. The intermediary then becomes responsible for forwarding the network data on to an appropriate base driver **104-108** which is then responsible for transfer of the data to the NIC hardware **116-120** for delivery over a network connection **122**.

On data reception over the network **122**, all NICs see the data, but only the NIC hardware with the appropriate matching MAC filter responds to the incoming data. If a NIC accepts network data, it is forwarded to its driver, which in turn forwards it to the intermediary layer which multiplexes the data to an appropriate protocol stack.

The intermediary layer is capable of accepting many upper-layer protocol stacks, in addition to multiple drivers below it. Although not provided by present networking environments, this ability provides an opportunity for allowing transparent fail-over, load-balancing, and support for new network protocols and features, without changing existing base drivers **104-108** for current network interfaces **116-120**.

In order to present a concrete example of how the intermediary layer **102** can be used to transparently extend the capabilities of existing network configurations, it will be assumed that a standard Novell Open Data-link Interface (ODI) based network is being extended. The ODI environment is a well known network configuration which does not provide for transparent logical replacement of a failed or failing NIC, nor support for features not natively supported by NIC networking software (e.g., the NIC drivers **104-108**). It will be clear to one skilled in the art how to implement the invention in other network environments, e.g., Windows 9x/NT, Macintosh, Unix / Linux, etc., as each environment uses an intermediary layer.

FIG. 2 shows the relationship between the traditional OSI model **200** and Novell's ODI **202**. The ODI is an application programming interface (API) developed by Novell for writing network drivers. As noted above, each supported NIC frame type is considered a different logical card, allowing the same NIC to carry data for different protocols (e.g., the same NIC can be simultaneously connected to both an IPX/SPX network as well as a TCP/IP network). As will be discussed, the ODI API can be extended to provide the claimed additional networking features.

Shown are the seven OSI layers **204-218**, and the corresponding ODI layers **220-226** that are different from the OSI model. The protocol stack (PS) **220** represents each protocol implementation (e.g., SPX, AppleTalk, etc.) available within a computer.

The Multiple Link Interface Drivers (MLIDs) **224** are device drivers which send and receive packets to and from the physical layer (or logical topology) **218, 226**. MLIDs have three portions, a Media Support Module (MSM), a Topology Specific Module (TSM), and a Hardware Specific Module (HSM). The HSM portion of MLIDs append or strip frame headers from data received over the physical layer, but they do not interpret packet data. Instead, received packets are passed on to the TSM and MSM for packet type classifying and then forwarded to Link Support Layer (LSL) **220** to be multiplexed to the appropriate protocol stack based on the contents of Event Control Blocks (ECBs).

ECBs are NetWare buffers that are used to send, receive and manage packet data. Generally an ECB contains information set by a protocol containing a block (e.g., 1500 bytes) of data. Such information includes the protocol originating the data, the NIC to which the data is to be sent, context information for the data, as well as other housekeeping information maintained by the protocol. The LSL effectively operates as a router, and coordinates communication between protocol stacks and MLIDs.

Traditionally, when data is transmitted, a protocol stack **220** receives data from an application program **204**. The PS **220** determines whether to split the data into fragments, and also determines the size of the fragments. The PS adds a protocol header to the data, destination MAC address and logical board through which packet will be transmitted are placed in the ECB, and sends the data to the LSL **222**. The LSL routes the data to an appropriate MLID **224**. The MLID adds a MAC header to the data,

and hands the data to the appropriate NIC (LAN adapter). Recall that the source MAC address added to the MAC header was read from the NIC during network initialization; this value is retained by the MLID for later use. (It is too time consuming to query each NIC for its MAC address each time data needed is transmitted or received – hence this determination is supposed to be only performed once.) NIC **226** adds a packet preamble and places the data on the wire (alternatively, the medium may be wireless).

Receipt of transmitted data operates conversely. On data receipt, the NIC strips the packet preamble from received data. Only packets which meet the NIC's stored MAC address passes the NIC's receive filter. The MLID strips the MAC header from the data, places various fields in an ECB and forwards the data pointed by the ECB to the LSL. The LSL routes the packet to the appropriate PS, and the PS removes the protocol header and transfers the data to the application.

Unfortunately, as noted above, in a traditional communication model, transmission and receipt breaks down if the NIC fails. Failure can happen by hardware failure, disconnected cable, or other communication breakdown. That is, since the data stream is hard-wired to travel to a particular NIC, a failed NIC will cause communication to fail.

FIG. 3 illustrates one implementation of the FIG. 1 communication configuration in the FIG. 2 Novell networking context. Shown is a protocol stack **300** corresponding to the protocol stack(s) **100** of FIG. 1, an LSL multiplexer **302**, and a virtual protocol stack (PS) **304** and virtual MLID **306**. The protocol stack is a protocol typically provided by the operating system vendor (or is a compatible protocol provided by a third-party, such as an network interface vendor which supports the network interface).

In a traditional communication configuration, the LSL routes communication data to MLIDs **308-310** according to MLID identifiers (usually a logical board number) embedded within such communication which correspond to stored within the PS, the LSL and the MLID for corresponding. Instead, here, the protocol stack **300** is bound to the virtual MLID **306**. Other protocols can be similarly bound. The virtual MLID appears to the LSL as a real adapter card's MLID driver. Consequently, network data **320** originating from the protocol stack **300** is routed by the LSL **302** to the virtual MLID **306**

as normal. As discussed above with respect to FIG. 1, after this data is received at the virtual MLID, it may be repackaged as if originating from a different protocol stack **304**.

This configuration allows the data **320** to be repackaged in a protocol format entirely different from the original protocol's format. Recall that the underlying philosophy to the OSI **200** / ODI **202** models is that each higher layer provides extra functionality and/or features to a communication stream, such as data integrity checks, reliability checks, out-of-order packet re-ordering, etc. Thus, the virtual protocol may re-transmit the data **320** in a specialized encryption format (e.g., IPsec), or in a virtual LAN format, or other format. (See also FIG. 5.)

Further, the repackaged data can be addressed so that it is directed towards any of the MLIDs **308-312**, allowing dynamic alteration of which MLID will be used for communication. The virtual MLID maintains a correspondence between originating protocol stack **300** for a particular data stream **320**, and the MLID **310** to which the data stream was delivered. When a transmit response is received by MLID **310**, the LSL **302** routes the response back to the "originating PS", i.e., virtual PS **304**. Internally, this response is stripped off any additional protocol features added to the outgoing data stream, and is presented back to the LSL **302** as if just received by the virtual MLID **306**. The LSL then routes this response back to the real originating protocol stack **300**, with the protocol stack **300** and LSL **302** being unaware of the data indirection.

One method for effecting the data repackaging is to copy and modify ECBs. In Novell networking, each transmitted data stream includes an Event Control Block (ECB) along with pointers to an ECB data payload (pointers within the ECB point to data buffers containing the payload). While redirecting the data **320**, the ECB for this data is copied into a second ECB (per networking conventions, the first ECB should not be modified). The second ECB's payload pointers (e.g., FragmentOffset, FragmentSize) remain directed at the first ECB's payload data buffers. To track the copy, the second ECB is adjusted to back-reference the original ECB. This can be accomplished through an entry in the copied ECB's ProtocolWorkspace portion. This second ECB is also adjusted to direct the data towards a MLID of interest. Similar header copying and adjustment may be used for other networking environments (e.g., Linux / Unix, Macintosh, Microsoft Windows 9x/NT, etc.). By creating copies of the ECBs, one can

control the format and destination of data originally received from the protocol stack
300.

Shown is a control line **322** (e.g., data pathway) in communication with each
MLID **308-312**. The control line **322** can be implemented using the Driver Management
5 Input/Output Control (IOCTL) method defined in ODI Lan Driver specification, where a
formatted management ECB is sent to a base driver. A command/query communication
data structure can be defined using a standard NetWare ECB buffer used in the Novell
driver management API. This command/query ECB will be filled with appropriate
command and related parameters, and then be handed to an MLID base driver by
10 calling the Ctl14_DriverManagement API. The base driver will act according to the
command field in this buffer and fill a response result and parameters in the same
buffer. In particular, an MLID can be directed to replace its stored MAC address with a
different value, thus replacing the value obtained when initializing the networking
hardware **330-334** and software (usually occurring at base driver load time). The MLID
15 stores the new MAC address in the NIC's address filtering hardware as well as in the
MLID's configuration table. Another care is taken to update the TSM's source MAC
address field it adds to every transmitted packet, with the new updated MAC address.

Once a MLID can be instructed to revise its stored MAC address, multiple NICs
can be installed in a machine, but where only one NIC **332** is actively communicating
20 through its MLID **310**, and the others **308/330, 312/334** are left "disabled" for fail-over
purposes. In one embodiment, an Adapter Fault Tolerance (AFT) mode is provided,
where a fail over NIC waits for failure of an active NIC. (See also FIG. 4.) If the active
NIC **332** suffers a hardware failure or its link is off, a fail over NIC **330** can be brought
on line by having its associated MLID **308** revise its stored MAC address **314** with the
25 MAC address **316** from the failed NIC.

Similarly, in one embodiment, load balancing may be effected by directing plural
MLIDs **308-312** to store an identical MAC address. The virtual PS **304** can then monitor
data traffic flow to identify a least-busy NIC and direct the LSL **302** to route data **320** to
that NIC. (The LSL does not route data to MLIDs based on the stored MAC address.)
30 In this embodiment, an Adaptive Load Balancing (ALB) mode, and Link Aggregation

mode is provided, where all NICs are active, until one fails, causing the failed NIC to be removed from load balancing until it is repaired.

Each NIC **330-334**, unaware of the MAC address stored in its associated MLID driver, receives the data stamped in the source address field by its MLID with the stored MAC address and submits it to the wire **336** for delivery.

Upon receipt of end of transmit response, the NIC presents the response to its MLID. The MLID passes the response to the LSL which routes it back to the virtual PS **304**. Recall that the virtual PS **304** and MLID **306** have maintained a concordance between data received from an originating protocol stack **300** and the MLID to which the data was ultimately forwarded. Based on this, the data can be returned to the originating protocol stack **300**.

FIG. 4 is a general flowchart for transparent fail-over for a defective NIC, such that networking software never becomes aware of the failure. When network data is received **400**, a test is performed to determine the status of the NIC. In one embodiment, this test is performed by polling **402** the base driver which in turn tests the network hardware if necessary. This polling and possible fail over may also occur asynchronously with network communication. As discussed above, in the Novell networking context, routing is performed by the LSL according to the contents of repackaged originating protocol stack data. A test **404** is performed to identify whether polling revealed a defective NIC. If a "primary" NIC fails the test, meaning that the NIC has or is in the process of failing, a fail over "secondary" NIC is identified **408**. As discussed above, one can have several secondary NICs installed within a computing device, but left idle until needed for fail over duty. One method for implementing such an idle state is to use the Novell Driver Management API to direct the MLIDs for the secondary NICs to enter into an inactive state. Then, when a primary NIC has failed, one of the secondary NICs can be selected and the corresponding MLID directed to load the MAC value for the failed NIC. The data from the originating protocol stack can then be directed to the fail over NIC until such time as the primary NIC is replaced or repaired.

The identified **408** secondary NIC is set **410** to utilize the MAC for the failed primary NIC. When asynchronously polling NICs, replacement of the failed NIC occurs before a sender of data receives a time out or other communication error from attempting communication through the failed NIC. In an alternate embodiment, NIC testing is performed along with data transmission, such that the base driver is polled **402** and evaluated **404** before data transferred **406** to a NIC.

FIG. 5 is a general flowchart for transparently re-formatting an original data stream from a protocol stack in a protocol format or networking feature not supported by the OS's protocol stack software. This figure shares items **400-410** of FIG. 4. This figure differs in that if **404** the NIC test yields okay, a subsequent test **412** is made to determine whether a new protocol or networking feature is to be used. For example, the NIC may be connected to a network requiring encrypted IP communication. If so, then the data can be repackaged **414** in an appropriate protocol and routed **406** to the NIC for distribution over a network.

Note that in an alternate embodiment, testing **402** a NIC, identifying **404** its health, and selecting a fail over NIC **408**, **410** can be performed asynchronously (e.g., in parallel) to regular network traffic processing. In this alternate embodiment, data is received **400** and processing continues directly with checking **412** whether a new protocol or feature is being used. A database, run-time variable, or other memory is used to store a MAC (or other identifier) address for a current primary adapter (e.g., the NIC presently in-use) to be used for network transfers. In parallel to processing received data, a separate process performs the testing **402**, **404** and fail-over selection **408**, **410** of a replacement NIC.

FIGS. 6 show API function entries for implementing the above-described on-the-fly modifications to networking communication. The API defines the control, management and advanced data interface between the invention, which includes a virtual protocol stack and virtual MLID as discussed above, and the MLID Base Driver in Novell's NetWare ODI networking environment. The API provides an interface for controlling and querying the base driver, as well as for transfer of data to and from a

base driver, and encapsulation of data in data formats (e.g., new protocols and new networking features) not currently supported by the standard Novell interface. As discussed above, underlying the API is the control channel (FIG. 3, item **322**), which among other actions, may be used to change a NIC's MAC address (Node Address) on the fly.

In an API configuration, the control interface **322** is based on a master slave protocol, where the virtual PS **304** and MLID **306** are the master, and the base drivers **308-312** are the slave. The slave drivers are polled at a predetermined polling period, e.g., **100** to **400** milliseconds or other period depending on data response requirements. As an alternative to polling, one can generate events from a base driver through the Novell Event Bus (NEB), with the virtual PS and MLID registered to receive events. However, for simplicity in presentation, it has been assumed that polling is utilized, and that the MLID base drivers respond to control messages immediately.

Starting on FIG. 6A, a first API command is IdentifyYourself **500**. This command requests a base driver (a MLID in a Novell networking context) to identify itself. This should be the first command sent to a base driver to allow verification that the base driver knows the virtual PS communication protocol. Passed as calling parameters to each API call are a VendorID **501**, and CommRevision **502**. The VendorID identifies the requestor to the base driver receiving the command, and allows the base driver to perform validation of the request. The CommRevision identifies the revision number of the API used to query the base driver, and allows the driver to ensure it responds appropriately to the API (e.g., different API versions may have different data expectations).

In response to the identity check, the base driver should return the following fields: CommRevision **503**, which indicates the driver's communication interface revision number; VendorID **504**, which indicates a vendor ID value of base driver; ResponseCode **506**, which embodies a response, if any, to commands sent to the base driver; and CopyrightString **508**, a protected statement returned by a driver to indicate that the driver is authorized to be communicating with the virtual PS. For example, to ensure that only licensed vendors are using the virtual PS to repackage network communication in a new (e.g., previously unsupported) format, the driver can be

required to return a copyrighted string before the virtual PS will communicate with that driver. Since the string is protected under copyright (and other applicable regulations), to be compatible with the virtual PS environment, one would have to seek proper authorization, or illicitly use the string. In this latter context, such illicit use may subject one to significant liability.

Another API command is ReportNodeAddress **510**. This command is used to identify a base driver's stored MAC address. As discussed above, this value can be determined by returned value from the driver. Or, alternatively, the address can be determined by inspecting the configuration table of the driver.

Another API command is UpdateNodeAddress **512**. This command is used to tell a driver to override its MAC address with a new MAC address **513** passed as a calling parameter. As discussed above, this command can be used at fail over time for fault tolerance, or when load balancing (e.g., link aggregation mode for aggregated bandwidth) where all underlying NICs use the same MAC address. As a result of this command, a base driver should update its MAC address in its configuration table, MSM (under Novell) shared data space, and in its hardware receive filtering table. If the driver (or its NIC) is not ECB aware, then the new address should be written on every packet the base driver delivers to the NIC (e.g., in the DriverSend routine). Writing the new address overcomes a limitation in the ODI specification and the way MSM and EtherTSM are implemented.

Note that fail over can be used to allow uninterrupted network communication during hot plug removal of a NIC, such as in a PCI environment. When a card is removed, the base driver will report (in response to the ReportStatus request) the removed card. If card was the primary adapter (e.g., in use), network communication will fail over to a secondary adapter with the secondary NIC getting the MAC address of the primary adapter. If it is desirable to maintain separate MAC addresses when the removed card is returned, the primary and secondary cards can be instructed to swap MAC addresses. The driver representing the card in hot plug event will be called to update its MAC address even though the NIC is not present. When the card is re-inserted, the driver should use the swapped address. This function can also be used in fail over conditions during Adaptive Load Balancing (ALB) (ALB balances outgoing

server traffic among multiple NICs, providing scalable bandwidth as well as automatic backup links through fault tolerance).

Another API command is ReportStatus **514**. This command is used to ask a base driver to report its running status. In response, a base driver can report operating conditions such as link change **516**, line speed **518**, duplex **520**, hardware failure condition **522**, hot plug PCI event (card removal) **524** or other results as desired. As discussed above, status will be polled at a predetermined interval (e.g., **100-400 400msec**) according to required response times.

Continuing on to FIG. 6B, another API command is ReportCapabilities **526**. This command is used to ask a base driver to report supported capabilities. Such capabilities can include support for ECB out of band information **528**, e.g., the use of an ECB to pass advanced data features on a per packet basis to and from a base driver, virtual LAN **530**, IP encryption **532**, etc. Note that some results, such as indicating virtual LAN capability **530**, necessarily imply other results, such as support for the out of band information **528**.

Another API command is ReportVlanCapabilities **534**. If virtual LAN support is present, this command can be used to ask a base driver to report what VLAN standards **536** it supports, whether VLAN mode can be enabled or disabled **538**, and whether the base driver supports VLAN RX filtering **540** (this support, if performed in hardware, will offload from a system's PCI bus any VLAN traffic not destined for the system).

Another API command is VlanControl **542**. This command is used to ask a base driver to change VLAN modes **544**, if changing modes is supported (depends on the supported standards **536**), or to set the VLAN filtering mode (if supported).

Another API command is SetControl **546**. This command is used to ask a base driver to revise basic functionality. This is an opened function and depends on functionality required of the base driver and/or NIC. This command can be used, for example, to enable or disable ECB out of band mode **548**. (At unbind time, original driver settings will be restored.)

Another API command is BD_Disconnect **550**. This command notifies a base driver that this new advanced control communication with the driver has been

terminated. This provides opportunity for the driver to recover, if necessary, or make adjustments due to the termination of communication.

FIG. 7 and the following discussion are intended to provide a brief, general description of a suitable computing environment in which the claimed invention may be practiced. The invention may be described by reference to different high-level program modules and/or low-level hardware contexts. Those skilled in the art will realize that program module references can be interchanged with low-level hardware instructions.

Program modules include procedures, functions, programs, components, data structures, and the like, that perform particular tasks or implement particular abstract data types. The modules may be incorporated into single and multi-processor computing systems, as well as hand-held devices and controllable consumer devices. It is understood that modules may be implemented on a single computing device, or processed over a distributed network environment, where modules can be located in both local and remote memory storage devices.

An exemplary system for implementing the invention includes a computing device **602** having system bus **604** for coupling together various components within the computing device. The system bus **604** may be any of several types of bus structure including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of conventional bus architectures such as PCI, AGP, VESA, Microchannel, ISA and EISA, to name a few. Note that only a single bus is illustrated, although plural buses typically achieve performance benefits. Typically, attached to the bus **602** are at least one processor **606**, memory **608**, storage device (e.g., fixed **610**, removable **612**, optical/laser **614**), video interface **616**, input/output interface port **618**, and network interface **620**.

The processor **606** may be any of various commercially available processors, including Intel processors, or the DEC Alpha, PowerPC, programmable gate arrays, signal processors, or the like. Dual, quad processors, and other multi-processor architectures also can be used. The system memory includes random access memory (RAM) **622**, and static or re-programmable read only memory (ROM) **624**. A basic input/output system (BIOS), stored in ROM, flash ROM, cached to RAM or the like,

contains routines for information transfer between device **602** components or device initialization.

The fixed storage **610** generally refers to hard drive and other semi-permanently attached media, whereas removable storage **612** generally refers to a device-bay into which removable media such as a floppy diskette is removably inserted. The optical/laser storage **614** include devices based on CD-ROM, DVD, or CD-RW technology, and are usually coupled to the system bus **604** through a device interface **626, 628, 630**. The storage systems and associated computer-readable media provide storage of data and executable instructions for the computing device **602**. Note that other storage options include magnetic cassettes, tapes, flash memory cards, memory sticks, digital video disks, and the like.

The exemplary computing device **602** can store and execute a number of program modules within the RAM **622**, ROM **624**, and storage devices **610, 612, 614**. Typical program modules include an operating system **632**, application programs **634** (e.g., a web browser or network application program), etc., and application data **636**. Program module or other system output can be processed by the video system **616** (e.g., a 2D and/or 3D graphics rendering device), which is coupled to the system bus **604** and an output device **638**. Typical output devices include monitors, flat-panels displays, liquid-crystal displays, and recording devices such as video-cassette recorders.

A user of the computing device **602** is typically a person interacting with the computing device through manipulation of an input device **640**. Common input devices include a keyboard, mouse, tablet, touch-sensitive surface, digital pen, joystick, microphone, game pad, satellite dish, etc. One can also provide input through manipulation of a virtual reality environment, or through processing the output from a data file or another computing device.

The computing device **602** is expected to operate in a networked environment using logical connections to one or more remote computing devices. One such remote computing device **642** may be a web server or other program module utilizing a network application protocol (e.g., HTTP, File Transfer Protocol (FTP), Gopher, Wide Area Information Server (WAIS)), a router, a peer device or other common network node, and

typically includes many or all of the elements discussed for the computing device **602**. To provide NIC fail over, the computing device **602** has plural network interfaces **620** (e.g., an Ethernet card) as described above, having updateable MAC address. These interfaces **620** are coupled to the system bus **604**, allowing communication with the remote device **642**. Both the local computing device **602** and the remote computing device **642** can be communicatively coupled to a network **644** such as a WAN, LAN, Gateway, Internet, or other public or private data-pathway. It will be appreciated that other communication links between the computing devices, such as through a modem **646** coupled to an interface port **618**, may also be used.

In accordance with the practices of persons skilled in the art of computer hardware and software programming, the present invention is described with reference to acts and symbolic representations of operations that are sometimes referred to as being computer-executed. It will be appreciated that the acts and symbolically represented operations include the manipulation by the processor **606** of electrical signals representing data bits which causes a resulting transformation or reduction of the electrical signal representation, and the maintenance of data bits at memory locations in the memory **608** and storage systems **610**, **612**, **614**, so as to reconfigure or otherwise alter the computer system's operation and/or processing of signals. The memory locations where data bits are maintained are physical locations having particular electrical, magnetic, or optical properties corresponding to the data bits.

Having described and illustrated the principles of the invention with reference to illustrated embodiments, it will be recognized that the illustrated embodiments can be modified in arrangement and detail without departing from such principles. For example, while the foregoing description focused -- for expository convenience -- on the Novell networking / ODI API environment, it will be recognized that the same techniques and analyses can be applied to different networking configurations, such as certain Macintosh, Microsoft (NTx Win9x and the like) and Unix environments. In addition, even though description or claim language may speak to only one or two network interfaces, it will be appreciated that the invention is applicable to devices simultaneously utilizing many network interfaces.

And, even though the foregoing discussion has focused on particular embodiments, it is understood that other configurations are contemplated. In particular, even though the expressions "in one embodiment" or "in another embodiment" are used herein, these phrases are meant to generally reference embodiment possibilities, and are not intended to limit the invention to those particular embodiment configurations. These terms may reference the same or different embodiments, and unless indicated otherwise, are combinable into aggregate embodiments.

Consequently, in view of the wide variety of possible networking environments, the detailed embodiments are intended to be illustrative only, and should not be taken as limiting the scope of the invention. Rather, what is claimed as the invention, is all such modifications as may come within the scope and spirit of the following claims and equivalents thereto.

What is claimed is:

1. An application programming interface (API) for enhancing data network communication, comprising:

an identify address function including programming instructions for identifying a stored node address stored by a base driver for a network interface associated with the base driver; and

an update node address function including programming instructions for directing the base driver to update the stored node address with a new node address in a configuration storage of the base driver, and in a receive address filtering table for the network interface.

2. The API of claim 1, wherein the identify address function includes submitting a request to the base driver, to which is received a response including the node address stored by the base driver.

3. The API of claim 1, wherein the identify address function includes programming instructions for inspecting the configuration storage of the base driver, such storage having an entry identifying the stored node address.

4. An API according to claim 1, further comprising:
a driver identification function including programming instructions for sending an identity-check request to the base driver, said driver providing a response selected from a group consisting of: a predetermined identifier, a base driver revision number, and an identification of a vendor of the base driver.

5. An API according to claim 4, wherein the predetermined identifier is a copyright string for the vendor of the base driver.

6. An article of manufacture, comprising a computer readable medium having encoded thereon programming instructions capable of directing a processor to perform operations of:

an identify address function for identifying a stored node address stored by a base driver for a network interface associated with the base driver; and

an update node address function for directing the base driver to update the stored node address with a new node address in a configuration storage of the base driver, and in a receive address filtering table for the network interface.

7. An API according to claim 1, further comprising:

a first transmission function including programming instructions for re-transmitting data, received in a compatible format from a network source, in an incompatible format to a network destination; and

a second transmission function including programming instructions for re-transmitting data, received in the incompatible format from the network destination, in the compatible format to the network source.

8. An API according to claim 7, further comprising:

a report capabilities function including programming instructions for sending the base driver a request to have the base driver report its capabilities;

a receive capabilities function including programming instructions for receiving a response including said capabilities;

wherein the incompatible format is formatted according to said capabilities.

9. An API according to claim 7, further comprising:

a virtual LAN function including programming instructions to direct the base driver to enter a desired virtual LAN operative state; and

a disconnect function including programming instructions to notify the base driver that the API has concluded communications with the base driver.

10. An article of manufacture, comprising a computer readable medium having encoded thereon instructions to direct a processor to perform operations of:

a first transmission function for re-transmitting data, received in a compatible format from a network source, in an incompatible format to a network destination; and

a second transmission function for re-transmitting data, received in the incompatible format from the network destination, in the compatible format to the network source.

5 11. An API according to claim 1 for providing transparent fail-over from a first network interface to a second network interface, further comprising:

a status function including programming instructions for polling a first base driver for the first network interface to detect a failure of said first network interface;

10 wherein the update node address function includes a function to direct a second base driver for the second network interface to store the node address of the first network interface as the stored node address for the second base driver.

12. An API according to claim 11, in which a Novell ODI compliant network is utilized for network communication, and wherein the update node address function uses at least one ODI MLID Control Routine.

13. An article of manufacture, comprising a computer readable medium having encoded thereon instructions to direct a processor to perform an API having:

an identify address function for identifying a stored node address stored by a base driver for a network interface associated with the base driver;

an update node address function for directing the base driver to update the stored node address with a new node address;

a status function in communication with a first base driver for the first network interface to detect a failure of the first network interface; and

25 a function to direct a second base driver for the second network interface to store the node address of the first network interface as the stored node address for the second base driver.

14. An API according to claim 1 for providing transparent load balancing of data transmissions directed towards the network interface by distributing such data across a second network interface, further comprising:

a queue monitoring function including programming instructions for detecting a workload of the first network interface; and

a distribution function including programming instructions for routing a portion of said data transmissions through the second network interface, said distribution function utilizing the update node address function to associate the node identifier of the first network interface with the second network interface.

15. A networking method, comprising:
receiving first network traffic with a protocol stack;
sending said first traffic to an intermediary layer;
routing said first traffic to a virtual interface driver;
repackaging said first traffic by the virtual interface driver, and providing said repackaged traffic to a virtual protocol stack;
sending said repackaged traffic to the intermediary layer; and
routing said repackaged traffic by the intermediary layer to an interface driver for the network interface.

16. A method according to claim 15, further comprising:
locating a fail over network interface having a node address memory;
identifying a failed network interface having a node address;
storing the node address in the node address memory; and
routing network traffic for the failed network interface through the fail over network interface.

17. A method according to claim 16, further comprising:
wherein said first network traffic is received in a first protocol format, and said repackaged traffic is in a second network protocol format different from the first protocol format.

18. A method according to claim 16, wherein locating the fail over network interface comprises:

submitting an node identification request to a base driver for a potential fail over network interface; and

receiving a response from said driver, said response including an authentication string;

5 verifying said authentication string has a predetermined value before said potential fail over network interface is used as the fail over network interface.

19. An article of manufacture, comprising a computer readable medium having encoded thereon instructions to direct a processor to perform the operations of:

10 receiving first network traffic with a protocol stack;

sending said first traffic to an intermediary layer;

routing said first traffic to a virtual interface driver;

repackaging said first traffic by the virtual interface driver, and providing said repackaged traffic to a virtual protocol stack;

15 sending said repackaged traffic to the intermediary layer; and

routing said repackaged traffic by the intermediary layer to an interface driver for the network interface.

20. A method for redundant networking in a network environment, comprising:
20 determining operative status of a first network interface having a first driver, and of a second network interface having a second driver with a driver memory for storing a MAC address for said second interface;

25 if the first network interface is inoperative, instructing the second driver to store the first network interface MAC address in the driver memory to allow processing by the second network interface of network traffic bound for the first network interface;

directing the second driver to activate the second network interface; and

directing the first driver to deactivate the first network interface.

21. A method according to claim 20, in which the network environment is a
30 Novell based network, and wherein ODI commands are issued to said first and second drivers.

22. A method according to claim 21, further comprising:

receiving first network traffic by a protocol stack;

forwarding said first network traffic to a LSL;

5 routing said first network traffic from the LSL to a virtual MLID, and deriving second network traffic from said first network traffic;

providing said second network traffic to a virtual protocol stack; and

forwarding said second network traffic to the LSL.

10 23. A method for load balancing transmissions over a network, comprising:

providing a first network interface having a first driver, said first interface having a communication load;

providing a second network interface having a second driver with an updateable driver memory for storing a MAC address;

15 receiving data from a network source;

apportioning said received data into a first and a second portion;

instructing the second driver to store the first network interface MAC address in the updateable driver memory;

20 sending the first portion to the first driver for network transmission; and

sending the second portion to the second driver for network transmission.

24. A method according to claim 23, further comprising:

wherein a communication protocol is bound to a virtual network interface for receiving the data from the network source, and wherein said apportioned data is re-
25 transmitted by a virtual protocol stack directing said apportioned data to the first and second drivers.

25. A method according to claim 24, wherein the virtual protocol stack repackages said apportioned data in a format unsupported by the network source.

30 26. A system, comprising:

means for identifying a stored node address stored by a base driver for a network interface associated with the base driver; and

means for directing the base driver to update the stored node address with a new node address.

5

27. A system according to claim 26, further comprising:

means for re-transmitting data, received in a first format from a network source, in a second format to a network destination; and

10 means for re-transmitting data, received in the second format from the network destination, in the first format to the network source.

42390.P7279

Apparatus and Method for
Networking Driver Protocol Enhancement

5

ABSTRACT

10

15

Application Programming Interface, methods and apparatus are disclosed for enhancing data network communication. In a network including a first and a second network interface, each interface has an associated MAC address, and each network interface has a driver storing the MAC address for its associated interface. Under certain circumstances, such as in a fail-over condition, or to improve throughput, the second driver is conditionally directed to replace its stored MAC address with the MAC address of the first network interface. Thus, the second network interface can process network traffic as if it were the first network interface. Disclosed are several features and advantages resulting from such MAC reassignment.

FIG. 1

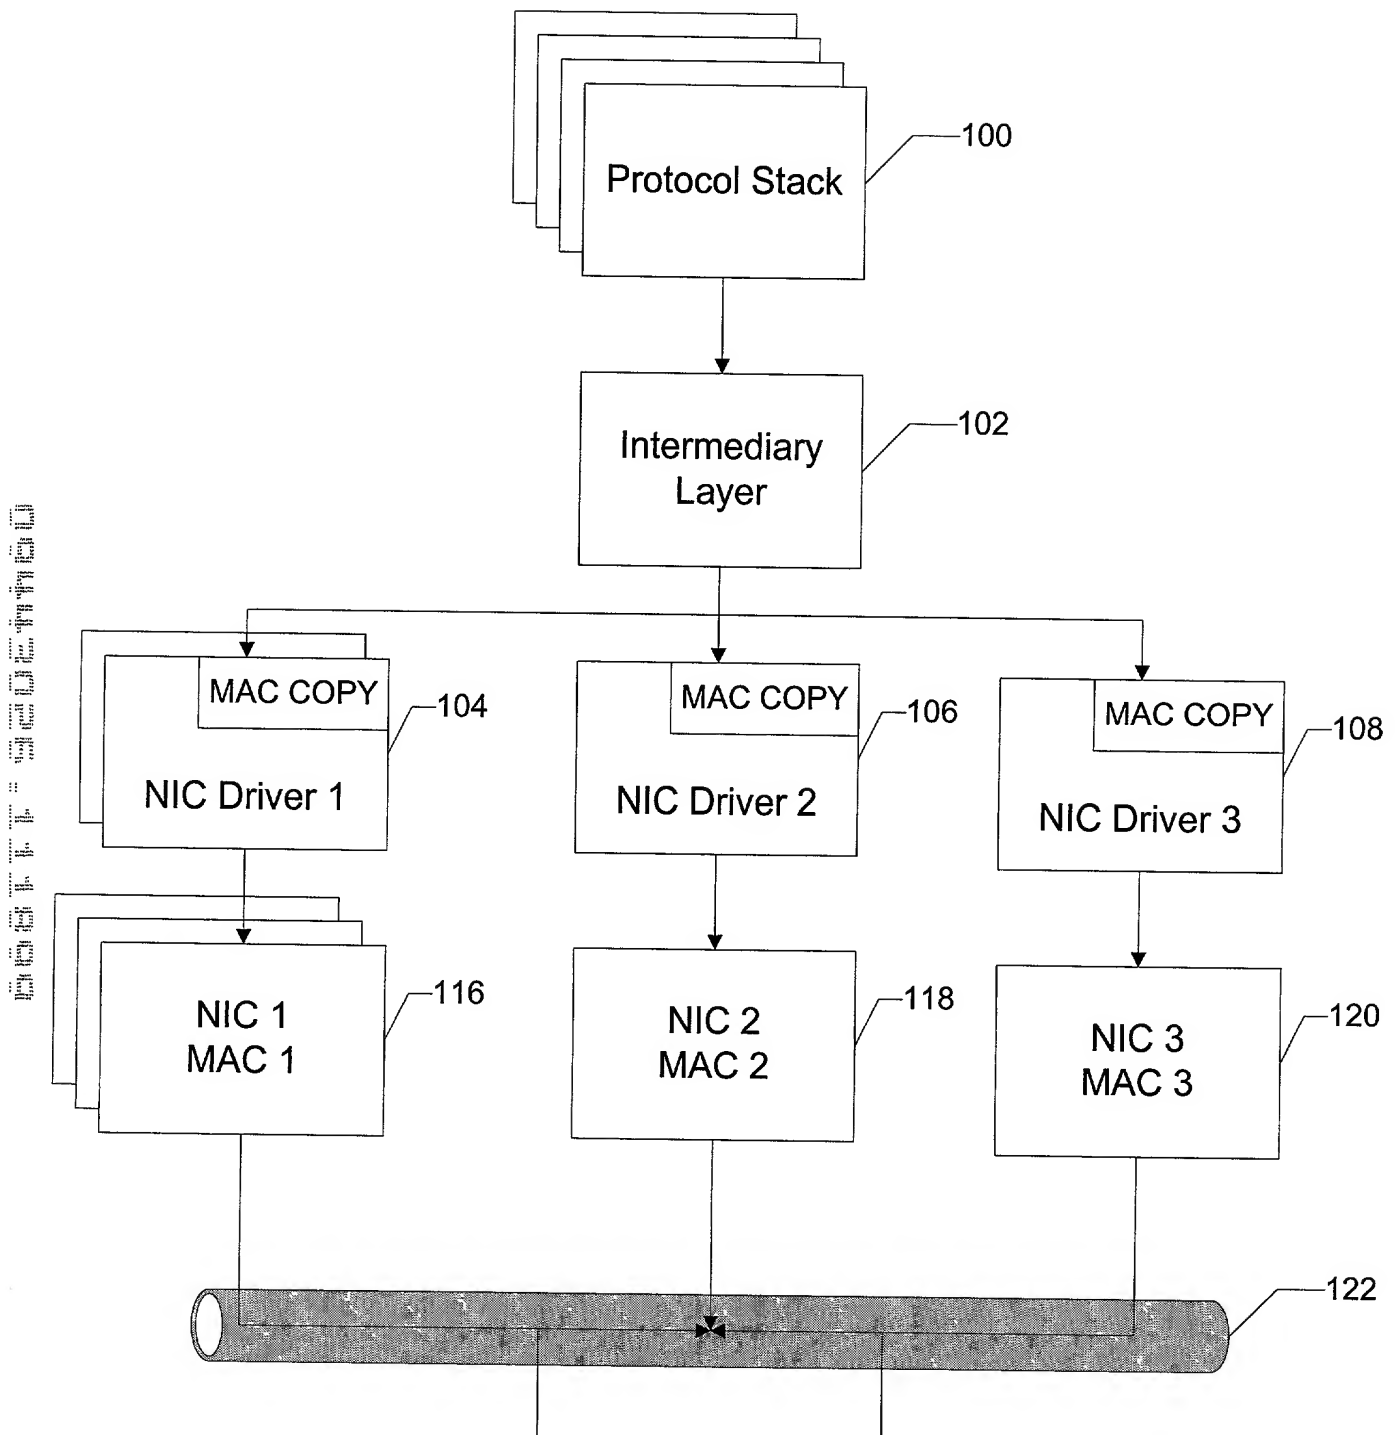


FIG. 2

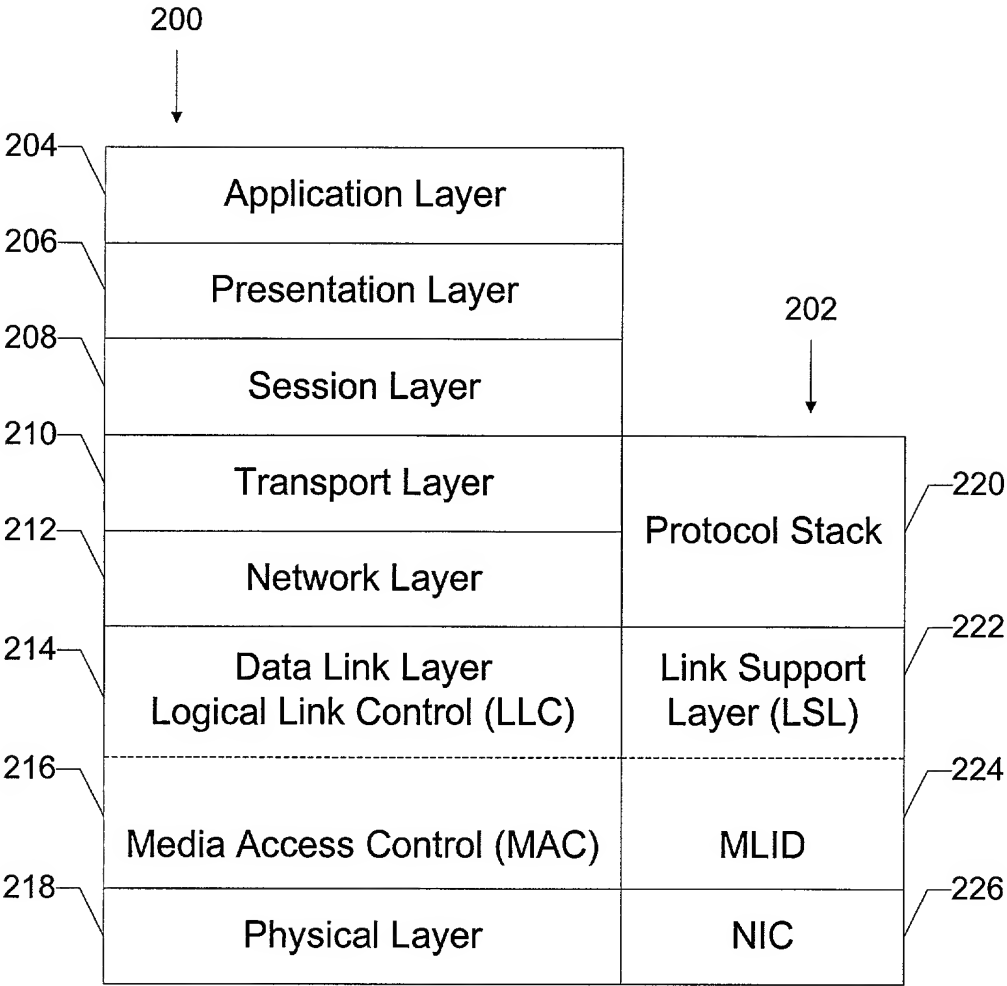


FIG. 3




FIG. 4

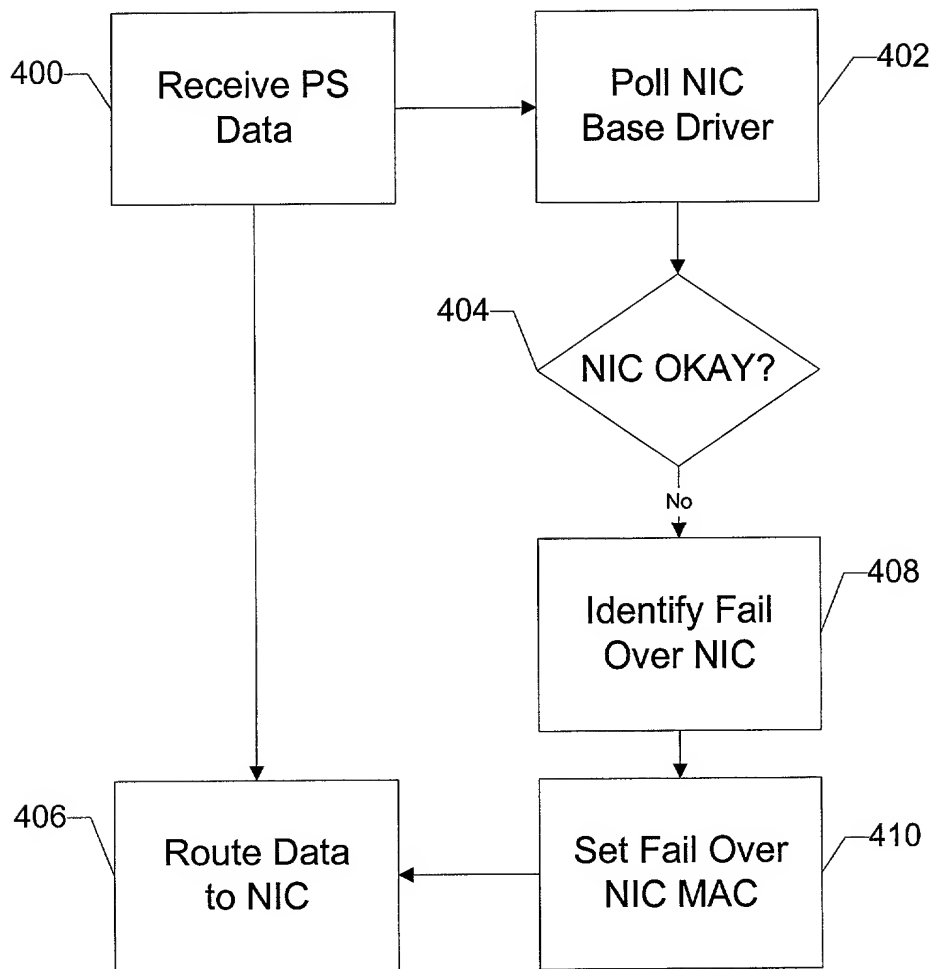


FIG. 5

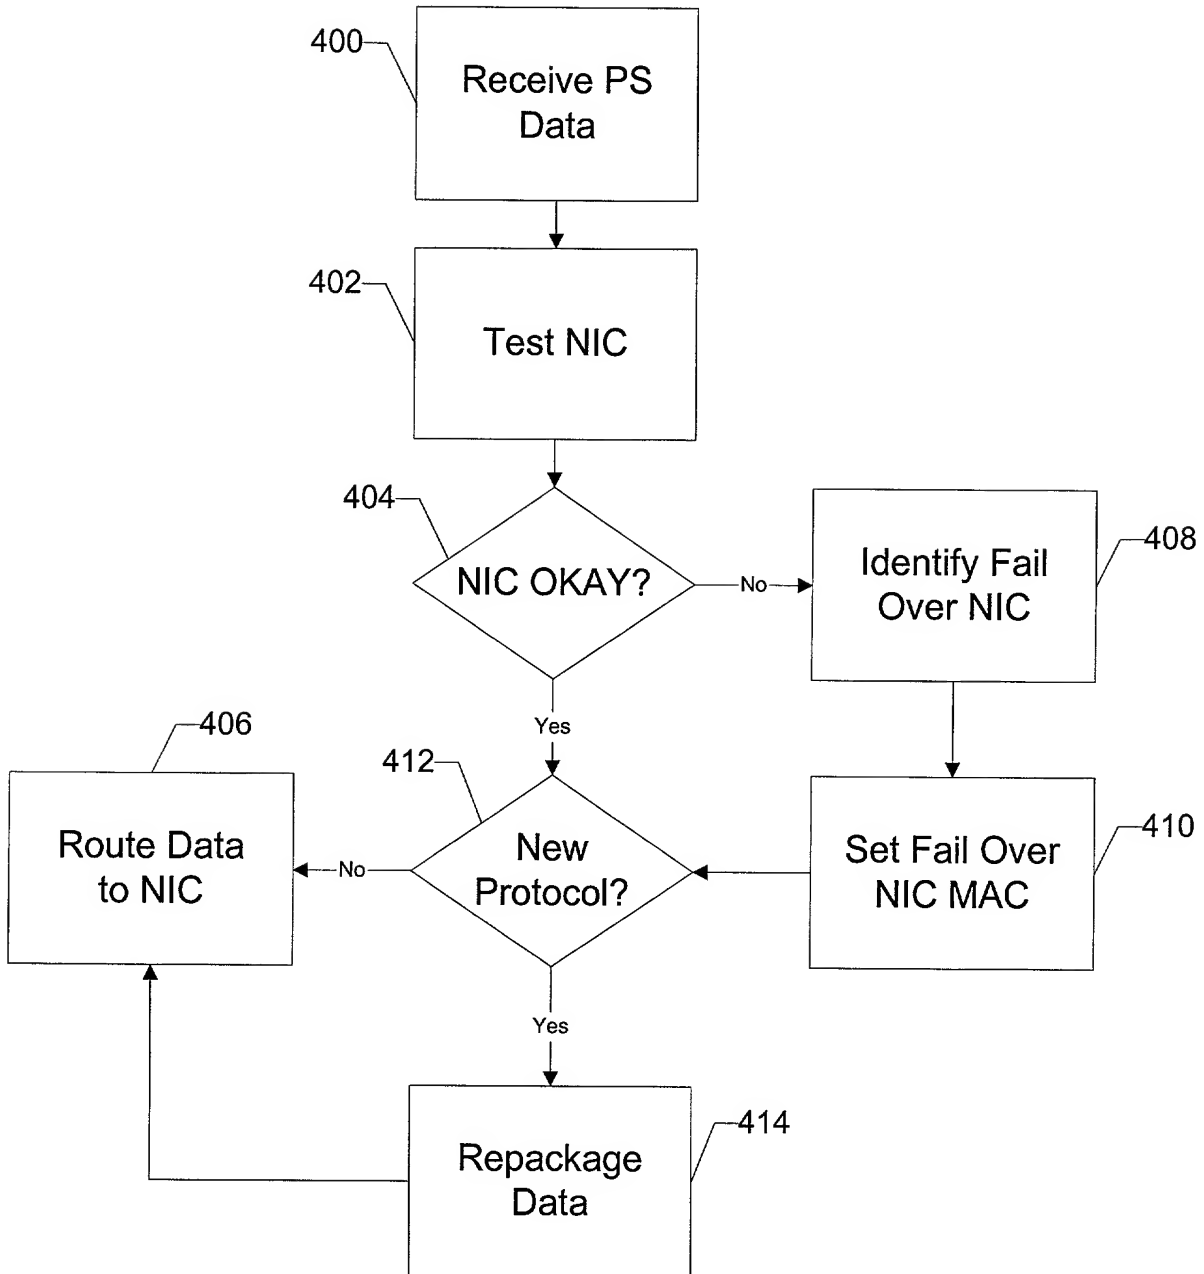


FIG. 6A

Function IdentifyYourself()

Arguments: Base Driver Identifier
VendorID: Vendor ID value of Intermediate driver
CommRevision: Intermediate driver communication interface revision number
Actions: Send control request to cause driver self-identification
Returns: CommRevision: driver communication interface revision number
VendorID: Vendor ID value of base driver
ResponseCode: Result code of any command to base driver.
CopyrightString: for verifying base driver authorization

Function ReportNodeAddress()

Arguments: Base Driver Identifier
VendorID: Vendor ID value of Intermediate driver
CommRevision: Intermediate driver communication interface revision number
Actions: Send control request to cause driver to report its MAC address.
Returns: MAC address

Function UpdateNodeAddress()

Arguments: Base Driver Identifier
VendorID: Vendor ID value of Intermediate driver
CommRevision: Intermediate driver communication interface revision number
MAC Address: New MAC address to be used by NIC.
Actions: Send control request to cause driver to replace stored MAC
Returns: Success/Failure

Function ReportStatus()

Arguments: Base Driver Identifier
VendorID: Vendor ID value of Intermediate driver
CommRevision: Intermediate driver communication interface revision number
Actions: Send control request to cause driver to report status
Returns: Link change
Line speed
Duplex
Hardware failure condition
PCI event

FIG. 6B

Function ReportCapabilities()

526 Arguments: Base Driver Identifier
 VendorID: Vendor ID value of Intermediate driver
 CommRevision: Intermediate driver communication interface revision number

Actions: Send control request to cause driver to report capabilities

Returns: Array containing identifiers for supported protocols, e.g.,

- 528 Out of Band
- 530 Virtual LAN
- 532 IP Encryption

Function ReportVlanCapabilities()

534 Arguments: Base Driver Identifier
 VendorID: Vendor ID value of Intermediate driver
 CommRevision: Intermediate driver communication interface revision number

Actions: Send control request to cause driver to report VLAN abilities

Returns: Array containing identifiers for supported VLAN protocols

- 536 Operating Mode
- 538 RX Filtering
- 540

Function VlanControl()

542 Arguments: Base Driver Identifier
 VendorID: Vendor ID value of Intermediate driver
 CommRevision: Intermediate driver communication interface revision number

544 VLAN mode

Actions: Send control request to cause driver to switch VLAN modes

Returns: Success/Failure

Function SetControl()

546 Arguments: Base Driver Identifier
 VendorID: Vendor ID value of Intermediate driver
 CommRevision: Intermediate driver communication interface revision number

548 Mode

Actions: Send control request to cause driver to switch to Mode

Returns: Success/Failure


Function Disconnect()

550 Arguments: Base Driver Identifier
 VendorID: Vendor ID value of Intermediate driver
 CommRevision: Intermediate driver communication interface revision number

Actions: Send control request to notify driver of communication end

Returns: Command receipt acknowledgment

FIG. 7



Attorney's Docket No.: 42390.P7279

PATENT

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION
(FOR INTEL CORPORATION PATENT APPLICATIONS)

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below, next to my name.

I believe I am the original, first, and sole inventor (if only one name is listed below) or an original, first, and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

**APPARATUS AND METHOD FOR NETWORKING DRIVER PROTOCOL
ENHANCEMENT**

the specification of which

XX is attached hereto.
_____ was filed on _____ as
United States Application Number _____
or PCT International Application Number _____
and was amended on _____.
(if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claim(s), as amended by any amendment referred to above. I do not know and do not believe that the claimed invention was ever known or used in the United States of America before my invention thereof, or patented or described in any printed publication in any country before my invention thereof or more than one year prior to this application, that the same was not in public use or on sale in the United States of America more than one year prior to this application, and that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representatives or assigns more than twelve months (for a utility patent application) or six months (for a design patent application) prior to this application.

I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d), of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

[illegible]

(Number)	(Country)	(Day/Month/Year Filed)	Yes	No
----------	-----------	------------------------	-----	----

(Number)	(Country)	(Day/Month/Year Filed)	Yes	No
----------	-----------	------------------------	-----	----

(Number)	(Country)	(Day/Month/Year Filed)	Yes	No
----------	-----------	------------------------	-----	----

(Application Number)	Filing Date
----------------------	-------------

(Application Number)	Filing Date
----------------------	-------------

(Application Number)	Filing Date	(Status -- patented, pending, abandoned)
----------------------	-------------	---

(Application Number)	Filing Date	(Status -- patented, pending, abandoned)
----------------------	-------------	---

I hereby appoint Farzad E. Amini, Reg. No. P42,261; Aloysius T. C. AuYeung, Reg. No. 35,432; Amy M. Armstrong, Reg. No. 42,265; William Thomas Babbitt, Reg. No. 39,591; Carol F. Barry, Reg. No. 41,600; Jordan Michael Becker, Reg. No. 39,602; Bradley J. Bereznak, Reg. No. 33,474; Michael A. Bernadicou, Reg. No. 35,934; Roger W. Blakely, Jr., Reg. No. 25,831; Gregory D. Caldwell, Reg. No. 39,926; Kent M. Chen, Reg. No. 39,630; Yong S. Choi, Reg. No. P43,324; Thomas M. Coester, Reg. No. 39,637; Roland B. Cortes, Reg. No. 39,152; Barbara Bokanov Courtney, Reg. No. 42,442; Michael Anthony DeSanctis, Reg. No. 39,957; Daniel M. De Vos, Reg. No. 37,813; Robert Andrew Diehl, Reg. No. 40,992; Tarek N. Fahmi, Reg. No. 41,402; James Y. Go, Reg. No. 40,621; Richard Leon Gregory, Jr., Reg. No. 42,607; Dinu Gruia, Reg. No. P42,996; David R. Halvorson, Reg. No. 33,395; Thomas A. Hassing, Reg. No. 36,159; Phuong-Quan Hoang, Reg. No. 41,839; Willmore F. Holbrow III, Reg. No. P41,845; George W. Hoover II, Reg. No. 32,992; Eric S. Hyman, Reg. No. 30,139; Dag H. Johansen, Reg. No. 36,172; William W. Kidd, Reg. No. 31,772; Michael J. Mallie, Reg. No. 36,591; Andre L. Marais, under 37 C.F.R. § 10.9(b); Paul A. Mendonsa, Reg. No. 42,879; Darren J. Milliken, Reg. No. 42,004; Thinh V. Nguyen, Reg. No. 42,034; Kimberley G. Nobles, Reg. No. 38,255; Babak Redjaian, Reg. No. 42,096; James H. Salter, Reg. No. 35,668; William W. Schaal, Reg. No. 39,018; James C. Scheller, Reg. No. 31,195; Anand Sethuraman, Reg. No. P43,351; Charles E. Shemwell, Reg. No. 40,171; Maria McCormack Sobrino, Reg. No. 31,639; Stanley W. Sokoloff, Reg. No. 25,128; Judith A. Szepesi, Reg. No. 39,393; Vincent P. Tassinari, Reg. No. 42,179; Edwin H. Taylor, Reg. No. 25,129; George G. C. Tseng, Reg. No. 41,355; Lester J. Vincent, Reg. No. 31,460; John Patrick Ward, Reg. No. 40,216; Stephen Warhola, Reg. No. 43,237; Charles T. J. Weigell, Steven D. Yates, Reg. No. 42,242; Reg. No. 43,398; Ben J. Yorks, Reg. No. 33,609; and Norman Zafman, Reg. No. 26,250; my attorneys, and James A. Henry, Reg. No. 41,064; Daniel E. Ovanezian, Reg. No. 41,236; Glenn E. Von Terssch, Reg. No. 41,364; and Chad R. Walsh, Reg. No. 43,235; my patent agents, of BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, with offices located at 12400 Wilshire Boulevard, 7th Floor, Los Angeles, California 90025, telephone (310) 207-3800, and Alan K. Aldous, Reg. No. 31,905; Robert D. Anderson, Reg. No. 33,826; Joseph R. Bond, Reg. No. 36,458; Richard C. Calderwood, Reg. No. 35,468; Cynthia Thomas Faatz, Reg. No. 39,973; Sean Fitzgerald, Reg. No. 32,027; Seth Z. Kalson, Reg. No. 40,670; David J. Kaplan, Reg. No. 41,105; Leo V. Novakoski, Reg. No. 37,198; Naomi Obinata, Reg. No. 39,320; Thomas C. Reynolds, Reg. No. 32,488; Steven P. Skabrat, Reg. No. 36,279; Howard A. Skaist, Reg. No. 36,008; Steven C. Stewart, Reg. No. 33,555; Raymond J. Werner, Reg. No. 34,752; Robert G. Winkle, Reg. No. 37,474; and Charles K. Young, Reg. No. 39,435; my patent attorneys, and Jeffrey S. Draeger, Reg. No. 41,000; Thomas Raleigh Lane, Reg. No. 42,781; Calvin E. Wells, Reg. No. P43,256; and Alexander Ulysses Witkowski, Reg. No. P43,280; my patent agents, of INTEL CORPORATION; and James R. Thein, Reg. No. 31,710, my patent attorney; with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

Send correspondence to Steven D. Yates, BLAKELY, SOKOLOFF, TAYLOR
(Name of Attorney or Agent)
& ZAFMAN LLP, 12400 Wilshire Boulevard 7th Floor, Los Angeles, California 90025
and direct telephone calls to Steven D. Yates, (503) 684-6200.
(Name of Attorney or Agent)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Sole/First Inventor Nimrod Diamant

Inventor's Signature _____ Date _____

Residence Kfar-Saba, Israel Citizenship _____
(City, State) (Country)

Post Office Address 4 Haegoz St. Kfar-Saba, Israel 44418
